

PATENT APPLICATION

**APPARATUS AND METHODS FOR CORRELATING MESSAGES SENT
BETWEEN SERVICES**

Inventor(s): Peter A. Panec
414 12th Street
Santa Monica, CA 90402
Citizen of United States of America

Jason Douglas
San Francisco, CA

Assignee: Grand Central Communications, Inc.

BEYER WEAVER & THOMAS, LLP
P.O. Box 778
Berkeley, CA 94704-0778
(510) 843-6200

PATENT

**APPARATUS AND METHODS FOR CORRELATING MESSAGES SENT
BETWEEN SERVICES**

By: Peter A. Panec and Jason Douglas

CROSS REFERENCE TO RELATED PATENT APPLICATION

This application claims priority and is a continuation-in-part of co-pending U.S. Patent Application, having Application No. 09/820,966, entitled "System and Method for Routing Messages between Applications", filed March 30, 2001 by Lev Brouk et al., which application is herein incorporated by reference in its entirety for all purposes.

BACKGROUND OF THE INVENTION

[001] The present invention relates to methods and apparatus for processing data within a computer network. More specifically, it relates to correlating messages sent within such computer network.

[002] Corporate reliance on technology has become more complex and more pervasive. Increasingly, companies are identifying opportunities to extend their core business or cut costs using the Internet. Both trends have put increasing priority on integrating disparate business applications. For this reason, enterprise application integration (EAI) has emerged as a solution for allowing information technology departments to build bridges that are designed to unify their legacy systems into a single enterprise application. Ideally, the creation of this single enterprise application would not require sweeping changes to the underlying structures.

[003] EAI suppliers typically offer end point solutions for managing business process interactions between end points within a computer network. Although a specific enterprise software package may be designed to transparently handle diverse business processes carried out by two or more end nodes, each specific enterprise software package requires releasing customized connectors or adapters which will work for the specific business processes and applications used by the specific end nodes. As a result, these enterprise solutions are not easily scalable. Additionally, scores of adapters need to be built for each vendor (e.g., Oracle, SAP and Peoplesoft). As each supplier releases new versions of their software, EAI vendors find themselves unable to gain traction under the burden of supporting their existing adapters.

[004] Notwithstanding the benefits of EAI, the software costs and resource investments of EAI prevent small-to-medium enterprise (SME) customers from embracing EAI solutions. For SMEs, reliance on application service providers (ASPs) represents an increasingly attractive alternative.

[005] The web service provider market is one of the fastest growing segments of the software industry. Service providers make enterprise applications (e.g., human resources administration, recruiting, travel and expense management, sales force automation) available to customers over the web at a server device. Those applications are fully managed and hosted by the provider, providing significant cost savings to enterprises.

[006] Some providers merely host and manage third-party packaged software for their customers ("managed hosters"). Others build new applications from the ground up to take advantage of the benefits and cost-savings of web service provider model. Service providers enjoy the profit margins and operational scalability of consumer Web companies like eBay

and Yahoo, while at the same time offering the feature sets of complex enterprise software applications such as Peoplesoft and Siebel.

[007] Although the service provider approach allows a single business to set up a host server for allowing itself and its business partners to use third party or customs applications, this approach does not allow the set up and dismantling of complex arrangements between business partners. For instance, a first business may wish to allow a second business to access a first set of services, while the second business may wish to provide a second different set of services to the first business.

[008] Additionally, message correlation in EAI or service provider approaches is nonexistent or cumbersome at best. In an end point based EAI, each end point must first agree on message identifiers before sending messages between each so that the related message can be correlated. This process is time consuming and difficult to implement and coordinate among diverse end point users.

[009] In view of the above, improved mechanisms for correlating messages sent between diverse entities in a scalable manner are needed.

SUMMARY OF THE INVENTION

[010] Accordingly, the present invention provides methods and apparatus for effectively correlating messages sent between services. In general, a message interchange network is configured to manage the use of services by remote entities or services within a computer network. In general terms, correlation information is stored for each message. The correlation information includes information for correlating each message with a specific call between two services and a specific session between two or more services.

[011] In one embodiment, a method for correlating messages sent between services within a computer network is disclosed. A message interchange network which manages a plurality of services which are each accessible by a plurality of services is provided. Correlation information regarding each message received into the message interchange network is tracked. The messages are generally being sent between pairs of the services, and the correlation information for each message pertains to each message and any other messages (e.g., along with their related services and routes) related to the each message.

[012] In one implementation, the correlation information for each message includes any of combination of (1) message information regarding the each message (2) call information regarding a call to which the each message and any other related message belongs, and (3) session information regarding a session to which the each message and any other related message belongs. In one aspect, the message information for each message includes a Hop Identifier (ID) uniquely identifying a hop between a sender and receiver of the each message and any combination of the following: an identification of the each message's sending and receiving service, an indication as to whether the each message has completed transmission, a reason or error log regarding why the each message has failed to complete its transmission

if the each message has failed, and a portion of the each message content, a size of the each message, a topic of the each message, a status on processing steps taken on the each message, and specification of any protocols used in receiving and sending the each message.

[013] In a further aspect, the call information for each call includes a Call Identifier (ID) uniquely identifying the each call and one or more of the following: an indication as to whether the each call is complete and a reason for the call not being complete if the each call fails to complete, a type of each call, a receiving and sending time for the each call, a sender and recipient service of each call, a status of policy evaluation for each call, and a set of hops in each call. In yet a further aspect, the session information for each session includes a Session Identifier (ID) uniquely identifying the each session and one or more of the following: an indication as to whether the each session is complete and a reason for the session not being complete if the each session fails to complete, a calculated or executed route for messages sent within the each session, and an identity and status of each service of the each session, an initiating time and completion time for each session, and an indication of a set of calls in each session.

[014] In one implementation, each message belongs to a particular call between two of the services. In a specific example, each call may include a request message and a response message or a notification message. In a general embodiment, a call is defined as a set of predefined message types, and a session is determined by the services which send messages for a set of calls.

[015] In one aspect, at least some of services are implemented on different computer systems and at least some of these computer systems differ from a computer system which implements the message interchange network.

[016] In a specific implementation, the tracking of correlating information includes (i) receiving a current message at the message interchange network, wherein the current message belong to a current session and a current call, (ii) when this is a first message received for the current session, assigning a session identifier for the current message and embedding the session identifier in the current message prior to forwarding it to its destination service, (iii) when this is a first message received for the current call, assigning a call identifier for the current message and embedding the call identifier in the current message prior to forwarding it to its destination service, (iv) assigning a hop identifier for the current message which uniquely identifies the current message, and (v) associating and storing the session identifier, the call identifier, and the hop identifier, along with message, call, and session information for the received message.

[017] In a further aspect, a query for correlation information regarding a particular session or call is received. The query is sent by a first one of the services. Correlation information is then sent to the first service related to the particular session or call of the query. In a further implementation, the correlation information includes information regarding messages sent between more than two services. In yet another aspect, it is determined whether the first service is authorized to make the query and the correlation information is only sent to the first service when it is determined that the first service is authorized.

[018] In another embodiment, the invention pertains to a computer system operable to correlate messages sent between services within a computer network. The computer system includes one or more processors and one or more memory. At least one of the memory and processors are adapted to provide at least some of the above described method operations. In yet a further embodiment, the invention pertains to a computer program product for

correlating message sent between services within a computer network. The computer program product has at least one computer readable medium and computer program instructions stored within at least one of the computer readable product configured to perform at least some of the above described method operations.

[019] These and other features and advantages of the present invention will be presented in more detail in the following specification of the invention and the accompanying figures which illustrate by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[020] The foregoing and other features and advantages of the invention will be apparent from the following, more particular description of a preferred embodiment of the invention, as illustrated in the accompanying drawings.

[021] FIG. 1 illustrates a message exchange network in accordance with one embodiment of the present invention.

[022] FIG. 2 illustrates components in a message exchange network in accordance with one embodiment of the present invention.

[023] FIG. 3 illustrates a request response pattern example.

[024] FIG. 4A is a diagrammatic representation of a communication exchange between a plurality of services in accordance with one embodiment of the present invention.

[025] FIG. 4B illustrates a group of related messages within another example communication exchange in accordance with one embodiment of the present invention.

[026] FIG. 4C is a flowchart illustrating a procedure for correlating messages in accordance with one embodiment of the present invention.

[027] FIG. 4D is illustrates a Message Correlation Table for correlating information regarding each message, a Call Correlation Table for correlating information regarding each call, and a Session Correlation Table for correlating information regarding each session, in accordance with one embodiment of the present invention.

[028] FIG. 5 is a diagrammatic representation of an example computer system in which embodiments of the present invention may be implemented.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

[029] Reference will now be made in detail to a specific embodiment of the invention. An example of this embodiment is illustrated in the accompanying drawings. While the invention will be described in conjunction with this specific embodiment, it will be understood that it is not intended to limit the invention to one embodiment. On the contrary, it is intended to cover alternatives, modifications, and equivalents as may be included within the spirit and scope of the invention as defined by the appended claims. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. The present invention may be practiced without some or all of these specific details. In other instances, well known process operations have not been described in detail in order not to unnecessarily obscure the present invention.

[030] In accordance with the present invention, the interchange of enterprise data is supported through an open platform for enterprise application integration (EAI). This open platform overlays a public network (e.g., the Internet) and does not require business entities to heavily invest in specialized software and resources. As will be described in greater detail below, the present invention enables the provision of extra-enterprise application integration as a service. This service facilitates EAI efficiently and affordably to the businesses that need it the most (i.e., the small- and medium-sized enterprise (SME) market). More generally, the open platform of the present invention can be used to support services provided by business-to-business (B2B) enablers, system integrators, and other node enablers.

[031] FIG. 1 illustrates a high-level overview of a message interchange system 100 according to one embodiment of the present invention. Message interchange system 100

includes a Message Interchange Network 150 that enables SMEs 110-m, webware ASPs 120-n, and in-transit processors (ITPs) 130-p to connect to one another, integrate business processes and applications, and exchange data for mission-critical business functions. In general, ITPs 130-p are operative to process messages that are in-transit from a sender to a recipient. ITPs 130-p can be designed to perform a variety of functions such as data transformation, enrichment, cross-reference ID mapping, filtering, credit scoring, or the like.

[032] A directory (not shown) includes a list of all SMEs 110-m, webware ASPs 120-n and ITPs 130-p that can be accessed via Message Interchange Network 150. Only publicly available services (i.e., those services that organizations register as accessible by any user of the network) are viewable in the directory.

[033] In general, all applications that are connected to Message Interchange Network 150 can be referred to as a service. In the illustrated embodiment of FIG. 1, applications owned by SMEs 110-m, webware ASPs 120-n and ITPs 130-p can each be referred to as services. Each service is owned by an organization, and an organization can have any number of services connected to Message Interface Network 150. The message exchange within Message Interface Network 150 is therefore between services.

[034] In one embodiment, services that receive messages can take a set of arguments that further define the intended action the service will perform on a received message. For example, a service may receive the name of an operation, or may permit configuration parameters. In this environment, the service would provide a means (e.g., through a URL to documentation) for message composers to know about arguments accepted by the particular service. The message composer can then include selected arguments as a part of the service declaration in a message.

[035] As described, services registered with Message Interface Network 150 represent applications that send or receive messages. An organization may, however, wish to create virtual services which act as proxies to other services. For example, a business X may have a relationship with business Y such that messages sent to business X's service are redirected to business Y's service. Services can implement redirection through routing scripts that map invocations of the service to invocations of another service, including redirection of replies.

[036] For each service registered by an organization with Message Interface Network 150, there are a number of properties and permissions that can be associated with the service. Examples include a unique service identifier, authentication information, mode of message delivery, windows of time during which messages are accepted, URL address of service, permission to invoke other services to act on a message, and rules that modify the invocation of services. These properties and permissions affect the routing of messages from or to the service.

[037] FIG. 2 illustrates the primary functional components that operate within message interchange system 100 in accordance with one embodiment of the present invention. The five primary functional components include service software development kit (SDK) component 210, web interface component 220, message router component 230, repository component 240, and billing component 250.

[038] SDK component 210 serves as a foundation for supported development of client applications that interface with Message Interface Network 150. Owning organizations can use SDK component 210 for custom integration with their software applications. As would be appreciated, SDK component 210 is not required to enable a service to access Message Interface Network 150. A service can use any development tool or process that would enable

the service to leverage the application programming interface (API) that is supported by message router component 230.

[039] In general, SDK component 210 enables the provision of an interface that would be available on most common platforms, and in most popular languages. In this manner, SDK component 210 abstracts away the complex technical requirements of transferring messages using Message Interface Network 150.

[040] SDK component 210 need not have any business logic built into it. SDK component 210 can be used to develop plug-ins to shrink-wrapped applications, thereby greatly reducing development time. As would be appreciated, SDK component 210 can provide convenient libraries and utilities that a service may optionally use to facilitate the (1) creation and reading of messages conforming to the message router component API, and (2) authentication of users of Message Interface Network 150.

[041] Repository component 240 is the primary database of Message Interface Network 150. Repository component 240 includes information on customer profiles, message logs, and directories. As will be described in greater detail below, message router component 230 uses repository component 240 to retrieve customer and application information that affects message routing. Message router component 230 also writes message log information to repository component 240 about messages that are processed through Message Interface Network 150.

[042] Billing component 250 uses the message log information in repository component 240 to derive actual usage of Message Interface Network 150 by customers, and handles the invoicing and payment collection from customers. In one embodiment, the billing within

message interchange system 100 can be based upon actual customer usage of Message Interface Network 150. For example, billing component 250 can charge customers based on a per transaction basis. In one embodiment, the per-transaction cost is based on the size of the messages being processed. As would be appreciated, these per transaction costs can be assessed to parties in a variety of ways. For example, the costs can be assessed against the originator of the message, the intermediate services, the recipient of the message, or any combination of those parties. This billing flexibility is in sharp contrast to conventional EAI solutions that generate revenue through software license fees.

[043] Web interface component 220 is the front-end component of Message Interface Network 150. Web interface component 220 interfaces directly with users by enabling login, registration, account maintenance, directory lookup, rules configuration, reporting, billing, and customer support functionality. The web interface provides online forms for data entry and can perform preliminary validations on the data. Through web interface component 220, the user can also perform queries against repository component 240 for directory lookups or reporting.

[044] In one implementation, Message Interface Network 150 is an open network architecture that not only facilitates the easy introduction of services into the network, but also enables businesses to access a robust suite of services via one connection to the Message Interface Network 150.

[045] As noted, message router component 230 provides the core function of message routing and delivery within Message Interface Network 150. In one embodiment, message router component 230 is implemented as an Internet-based message exchange service that

provides a transport level messaging service. In other words, message router component 230 need not be aware of the application semantics of a message exchange.

[046] Preferably, Message Interface Network 150 need not inherently provide business process modeling. This is in contrast to conventional EAI solutions that may require a continual traversal up and down a protocol stack in routing a message from a sending service to a recipient service. For example, if the protocol stack included transport, routing, transformation, and work flow layers, then each message exchange segment may require analysis and processing at each layer to determine the next service (intermediate or final) that should receive the message.

[047] As noted, services can post messages to and retrieve messages from message router component 230 using an API. This provision of a standardized interface enables parties to easily connect to and use Message Interface Network 150 without being restricted in the type of message content.

[048] In one embodiment, the protocol for posting and retrieving messages with Message Interface Network 150 is the Simple Object Access Protocol (SOAP). The SOAP messaging protocol defines a mechanism to pass commands and parameters between HTTP clients and servers. Through this standard object invocation protocol, HTTP is used for transport and XML is used for data encoding. The SOAP messaging protocol does not rely on the use of particular operating systems, programming languages, or object models on either the server side or the client side. As would be appreciated, other protocols can also be supported by Message Interface Network 150.

[049] While the message header can use extensible markup language (XML) syntax, the message body can accommodate any type of data, whether it be text or binary, encrypted or unencrypted. If the message body is also in XML form, then the message body can opt to use a schema based on an industry standard such as ebXML, BizTalk, RosettaNet, OAGIS, or any other suitable standard.

[050] In one embodiment, message exchange through Message Interface Network 150 is asynchronous. Recipient services can be configured to poll Message Interface Network 150 for incoming messages, or, if they have their own server, can have Message Interface Network 150 push messages to them.

[051] After a sending service posts a message to Message Interface Network 150, one or more in-transit services 130-p can operate on the message before it reaches the recipient service. In-transit services can perform useful operations on messages, such as data transformation, enrichment, cross-reference ID mapping, filtering, credit scoring, or the like. Through the standardized interface, in-transit services 130-p can independently join the Message Interface Network 150 and operate on messages. This flexibility encourages independent third parties to build services that can be plugged into Message Interface Network 150. Such an open network would encourage third parties to market a data service that generates revenue based upon the level of utilization of the service.

[052] As noted, in-transit services can be included in a message path that begins at a sending service and terminates at a recipient service. As will be described in greater detail below, sending services can explicitly specify a set of services to operate on a given message. In addition, recipient services can specify services that should operate on messages

before delivery to the recipient service. In one example, a recipient may always want messages to pass through a filtering service to screen out messages from unknown senders.

[053] Messaging through Message Interface Network 150 can be as secure as the participants desire. Each service registered with Message Interface Network 150 can specify a security policy declaring encryption and authentication levels for Message Interface Network 150 to enforce. For messages that flow through in-transit services, a sender can also specify the permissions for each in-transit service to access or operate on parts of the message.

[054] In one embodiment, Message Interface Network 150 uses the secure HTTPS protocol to support secure transport connections when a service posts a message or polls for messages, and when Message Interface Network 150 pushes messages to a client server. Authentication can be based on either username/password or certificates.

[055] SSL encryption as part of HTTP can be used to provide data protection during message transmission over the public Internet. In general, this level of protection is sufficient for most situations. Services can, however, perform their own extra encryption of message documents to keep them private even from Message Interface Network 150. Services that add extra encryption should ensure, however, that all services that operate on the message documents have the necessary keys to decrypt the documents.

[056] As is well known, the authentication protocol of SSL includes a server's presentation of a certificate to clients. Accordingly, Message Interface Network 150 presents a server certificate to services that connect for posting or polling. The connecting service has the option of then providing either a username/password or certificate for Message Interface

Network 150 to authenticate the service. The form of client authentication is a configuration option within the profile Message Interface Network 150 maintains for each service.

[057] When Message Interface Network 150 pushes messages to a service, the service's server should present a server certificate to Message Interface Network 150 for authentication of the service. For the reverse authentication, the service can then require either a username/password or certificate from Message Interface Network 150. Again, that option can be configured in the profile information Message Interface Network 150 maintains for the service.

[058] As a message flows through a selection of services on the way to the recipient service, and as the recipient service's response returns to the sending service, Message Interface Network 150 maintains an audit trail of all operations on the message and all services that touched the message. The audit trail serves several purposes. First, it enables Message Interface Network 150 to reconstruct the message history in the case of queries on the message trail. Second, it allows Message Interface Network 150 to compile a usage report for any service for reporting and billing purposes.

[059] Having described the general framework of Message Interface Network 150, a more detailed description of a message transaction lifecycle within Message Interface Network 150 is provided with reference to FIG. 3 in accordance with one embodiment of the present invention.

[060] In this framework, a message can be embodied as a self-contained collection of information to serve a particular purpose, such as a request, a response, a notification, or an acknowledgement. As noted, Message Interface Network 150 can generally be agnostic

about the content of a message other than header information that affects routing of the message.

[061] In one embodiment, request, response, and notification messages can be defined. A request message expects a subsequent response message from the recipient(s) to be returned to the sender. Request messages may represent inquiries, but might also represent update requests that only expect a return status response. If an error occurs in routing a request message, Message Interface Network 150 returns an error response message to the sender.

[062] A response message is issued by a recipient of a request message. The response message references the original request message. Failure of the response message may result in an error response message being returned to the sender of the original request message.

[063] A notification message is a one-way message. No response to the notification message is expected back to the sender. Message interchange network 150 can regard any response message referencing a notification message as an invalid message. If a notification message fails, no error message is returned to the sender.

[064] As would be appreciated, further messages can be defined for Message Interface Network 150. For example, a cancel message can also be defined, wherein the cancel message is used by the sender to cancel a previous message.

[065] The operation of these messages is now described with reference to the request/response illustration of FIG. 3. This illustration demonstrates a typical example of a sending service 310, such as an enterprise, making an inquiry to a recipient service 360, such as a webware provider. In one embodiment, a sender's application 312 that connects to Message Interface Network 150 is a desktop application. In another embodiment, a sender's

application 312 that connects to Message Interface Network 150 is an enterprise server, or an EAI package.

[066] The first step in the message transaction process is the creation of a message. In one embodiment, a sender formats the messages to conform to an XML schema for messages. This XML schema prescribes the format for message headers while allowing any kind of data to be included in the message body (or payload). As part of message construction, sending service 310 specifies the recipient service(s) 360 of the message. In one embodiment, a recipient service's name includes an organization and a specific service provided by that organization. The service name can be generally represented in the message via a globally unique ID.

[067] The actual set of elements contained in a message depend on whether the message is being posted or delivered. In one embodiment, a message includes a header element, a body element and/or attachments. In one embodiment, the attachments are based on multi-part Multipurpose Internet Mail Extensions (MIME).

[068] An embodiment of a message includes header and body elements. As would be appreciated, the actual message format can differ depending on the protocol. In particular, protocols other than the SOAP protocol can be used. The header element may include routing and invocation information. The header as posted by a sending service may often be modified by Message Interface Network 150 for delivery to the receiving service.

[069] The body element may include the documents the sender is sending to the recipient(s). These documents can also be operated upon by one or more services. The documents can be in the form of XML or any other representation, including text, binary,

etc. In one embodiment, all or part of the documents that are being sent are included in an attachment to the message.

[070] While messages preferably have a similar overall structure, the actual composition of elements can differ between the various message types and between messages as posted and as delivered. For example, some elements in a sent message can be changed or not be included in the message as delivered, such as elements particular to constructing a route. Some elements can also be inserted only in the message as delivered, such as identifier elements.

[071] If the sending service wishes to have the message routed through any services before delivery to the recipient service(s), the sending service can specify an explicit sequence of services that should operate on the message. The sender can also implicitly include services in the route for a message through the specification of routing scripts associated with the defining service. Routing scripts are described in greater detail below.

[072] After a message is constructed, the message is posted to Message Interface Network 150. FIG. 3 illustrates an example posting of a message by application 312 to message post interface 324. As noted, in one embodiment, the posting of a message is performed using the SOAP messaging protocol.

[073] If sending service 310 posts a message that does not have well-formed XML, the message posting is rejected and an error response is returned. In general, messages can be rejected for a variety of other reasons. For example, a message can be rejected if the service indicated in the message header as the sender is not the same as the actual sender of the message, the message is a duplicate posting of a previous message, a service attempts to

reply to a message for which it was not a recipient, or a response message does not reference a prior message.

[074] In one embodiment, each message posted by a service can have a unique handle assigned by the service to identify the message. This unique handle can be used to provide a means for Message Interface Network 150 to detect duplicate postings of the same message. Duplicate postings can occur in the case of failure recovery by the service. In one embodiment, if a service desires that Message Interface Network 150 should reject duplicate postings of a message, then the service could provide unique handles for messages and set a "potential duplicate" flag in messages that may be a duplicate posting. It should be noted that regardless of whether or not a service provides a unique handle for a message, Message Interface Network 150 can assign a globally unique session identifier to each posted message.

[075] After a message is posted, Message Interface Network 150 routes the message to the recipient service(s) 360. The routing of the message is based upon a route calculation by Message Interface Network 150. The calculated route includes all intermediary services 350 that are scheduled to operate on the message en route to recipient service(s) 360. The calculated route can be based on routing instructions specified explicitly in the message header and/or on routing scripts pre-defined by the sending service 310, recipient service 360, or any in-transit services 350 that have been included within the calculated route.

[076] In general, routing scripts define a procedure for enabling determination of at least part of a route. This procedure can be based on any type of criteria. For example, a procedure can be defined that determines a next destination of a message based on the existence of one or more attributes of the message. In another example, a procedure can be defined that

effects a determination based on the comparison of one or more attributes of the message to a reference value. In yet another example, a procedure can be defined that effects a determination based on pattern matching (e.g., regular expression matching). As would be appreciated, routing scripts can embody any of a variety of criteria-based procedures.

[077] Routing scripts can specify a sequence of services that operate on either inbound or outbound messages for a service. As noted, in-transit services may themselves have routing scripts requiring processing by other services. Therefore, the route calculation can be recursively defined based upon routing scripts specified by all services that interact with the message.

[078] In one example, the sending service 310 may specify a routing script that requires a cross-reference mapping service to be included in the calculated route whenever sending service 310 sends a message to recipient service 360. In another example, recipient service 360 may specify a routing script that requires that any incoming request messages must first pass through a filter service to block messages from a list of sending services 310.

[079] Routing scripts enable sending services to include services into the message route without having to explicitly specify the services in the message itself. Also, routing scripts enable recipient services 360 to require services 350 to be in the calculated route, regardless of the sending service's route specification.

[080] In one embodiment, a routing script is embodied as a routing rule. A routing rule includes two parts: a condition and one or more resultant actions. The conditional part of a rule can be based on any elements or element attributes in a message's header. Additionally,

content-based routing can be supported through conditional rules based on attributes of an element in a message's body and/or attachments.

[081] Every rule should have at least one condition. Conditions include an operator and zero or more operands. Example operators include equals, notEquals equalsOneOf, lessThan, greaterThan, and exists operators. In one embodiment, operators act on XML elements, XML attributes, or on other conditions.

[082] From the standpoint of the element operators, XML elements contain either child elements or character data. Therefore, the operands for an element comparison both represent the same type of content: either elements or character data. Character data can be in the form of a string, number, or date. Conditions involving elements that do not appear in the message will evaluate to false.

[083] Attributes generally have a type of character data, which can be string, number, or date. Many attributes are implicitly included in an XML document with default values. Therefore, an attribute identified in a condition can refer to either an explicit or implicit attribute. Conditions involving optional attributes that do not appear in the message will evaluate to false.

[084] The usual boolean operators can combine conditions into more complex conditions. Condition operators act on other conditions. Example condition operators include AND, OR, XOR, and NOT condition operators.

[085] The result of satisfying a rule's conditions is that an action will be triggered to modify the route for a message. Probably the most common result of a rule is to add one or more

services into the route for a message. Several rule actions can be defined, and several are described further in the above reference Lev Brouk et al. U.S. patent application.

[086] A service should maintain an evaluation sequence for the scripts associated with each role that the service can have with respect to a message. That sequence determines the order in which the scripts for that service are applied.

[087] In one embodiment, scripts are evaluated in the following order: (1) scripts for the sender of the message; (2) scripts for services included by the scripts for the sender (this is recursive); (3) scripts for the recipients of the message, in the order of recipients in the message header; and (4) scripts for services included by scripts for the recipients (this is recursive).

[088] When multiple scripts for a service include services into a route, the order of services in the route will follow the order of the scripts. That is, if script 1 inserts service A, and script 2 inserts service B, and if script 1 is evaluated before script 2, then service B follows service A in the route.

[089] In one embodiment, routing scripts are evaluated only once during the initial calculation of the route for a message. The message header contains the basic information to initially construct a route, such as sending service 310 and recipient services 360. The message can also contain an explicit specification of a set of services to include in the route. Once the route is constructed from the header information, routing scripts are applied to further elaborate the route.

[090] In an alternative embodiment, at least part of the message route is calculated after the physical routing of the message has begun. Dynamic routing is described in greater detail below in the context of physical and logical routing.

[091] At the transport level, Message Interface Network 150 routes a message to a service by delivering the message through the Internet to a physical machine on which the service resides. That service operates on the message and, if the message is a request, returns a response message back through the Internet to Message Interface Network 150. The sequence of message deliveries and responses between Message Interface Network 150 and services represents the physical routing of a message.

[092] Message interchange network 150 also provides a mechanism for a service to act on a message without the message being physically delivered to the service over the Internet. This mechanism is enabled through the logical routing of the message to the service. With logical routing, a service can modify the routing of the message or modify the context of the message for delivery to the next service. Significantly, a service can be logically included in a message routing, without being included as part of the physical routing of the message.

[093] In one embodiment, logical routing of messages is implemented through the specification of routing scripts. As described above, a service can define one or more routing scripts. These defined routing scripts are stored within Message Interface Network 150 and are processed to determine what routing behavior should occur when a message is logically routed to the service.

[094] Logical routing can take place statically or dynamically. With static logical routing, a message is logically routed to all services prior to any physical routing. In other words,

Message Interface Network 150 logically routes the message to all services prior to the physical delivery of a message to any services. This logical routing is represented by the sequential evaluation of the routing scripts that are defined by those services. As noted above, in one embodiment, the routing scripts are evaluated in the following order: (1) scripts for the sender, (2) scripts for the services included by the sender (recursive), (3) scripts for the recipient, and (4) scripts for the services included by the recipient (recursive).

[095] In dynamic logical routing, the logical routing is not completed prior to the start of the physical routing. Rather, the logical routing takes place in sequence with the physical routing of the message. The relation between logical routing and physical routing is described in greater detail below.

[096] As noted, Message Interface Network 150 delivers a message logically to every service participating in a message's routing. Of those services, some subset will also accept physical delivery of the message.

[097] To illustrate this concept, consider an example where service A includes service B into the message route. Service A can include service B into the route either prior to itself in the route (provided service A is not the originator of the message) or after itself in the route. In either case, Message Interface Network 150 would logically route the message first to service A, which includes service B into the route. Message interchange network 150 then logically routes the message to service B, and after service B produces a response, Message Interface Network 150 logically returns the response to service A. The point at which service A physically receives the message depends on whether service A included service B prior or after itself in the route. If service A includes service B prior to itself in the route, then the order of physical delivery is first to service B then to service A. Conversely, if service A

includes service B after itself into the route, then the order of physical delivery is first to service A then to service B. In the latter case, the response from B is not necessarily physically delivered back to service A. Rather, it may be only logically delivered back to service A.

[098] Services to which a message is logically routed do not necessarily have to also physically receive the message. In the above example, service A could have been logically routed, with physical delivery only to service B. Consider the following scenario. Suppose service X includes service A into the route and service A includes service B into the route. The logical routing of the message would proceed from service X to service A to service B back to service A back to service X. Service A can choose not to be included into the route for physical delivery, in which case the physical routing of the message is from service X to service B.

[099] In general, the act of routing a message (physically or logically) to a service can be thought of as an invocation of the service. When a service includes another service into the route of a message, the including service is effectively invoking the included service. The invocation of a service does not necessarily imply the physical delivery of information to the invoked service. The logical routing of a message is then the logical invocation of services. A route that includes a progression of services including other services can effectively be modeled as a progression of invocations.

[0100] In logical routing, each service is not only able to manage the inclusion of other services into the route but is also able to manage the context of those inclusions. From the standpoint of invocations, an invoking service is able to set the context for the invocation.

An invoked service can also set the context of its return.

[0101] Message interchange network 150 can also affect context management on behalf of invoking services. As noted, while an invoking service can be logically included in a message routing, it need not be included as part of the physical routing of the message. In general, Message Interface Network 150 can persistently store contexts of a message, thereby enabling proper restoration of contexts upon return from an invocation. Several embodiments of context management are described in the above referenced Lev Brouk et al. U.S. patent application.

[0102] Having touched on a framework for logical routing and invocation of services, the description of the physical routing process is continued with reference to FIG. 3. In this example, it is assumed that static logical routing has produced a route for the message.

[0103] After deriving the route for a message, Message Interface Network 150 validates the route. There are numerous conditions that can cause a route to be invalid. For example, there may be routing permission violations or a service may be currently disallowed by Message Interface Network 150 due to, for example, the non-payment of usage bills.

[0104] If the route is determined to be invalid, then Message Interface Network 150 rejects the posted message and may return an error to the sending service 310 either in the response to the posting call or in an error message.

[0105] Message interchange network 150 routes the message to all services in the calculated route. In the event of failure at any stage in the routing, Message Interface Network 150 aborts the message routing and, if the original message was a request, returns an error message back to sending service 310. For messages with multiple recipients, an error in the routing to one recipient will not necessarily affect routing to other recipients. Errors during

routing can occur due to several circumstances. For example, a message may fail to reach a recipient within the expiration time for the message, a service may fail to return a reply to a delivered message, or a service may return an error status for a message.

[0106] Message Interface Network150 sequentially delivers a message to each service identified in the message route. In FIG. 3, this process is illustrated as a flow of the posted message through message routing element 340, and on to one or more services 350. In one embodiment, Message Interface Network150 includes all of the message documents in the message delivered to service 350, even if service 350 only expects to operate on one document or documents of a particular content-type. Service 350 would ignore documents that it does not expect.

[0107] As noted, Message Interface Network150 invokes in-transit services in the same way as delivering a message to any sending or recipient service. In general, a service does not necessarily need to be aware whether it is being invoked as an in-transit service or as a recipient service.

[0108] After processing the message, service 350 sends the results in a response message back to Message Interface Network150. If service 350 is unable to produce a valid result for its operation on a message, then service 350 may return an error code in its response message to Message Interface Network150. If service 350 fails to respond to a received message, the message will ultimately expire, and Message Interface Network150 may return an error message back to the sending service 310.

[0109] Upon receipt of the response message from service 350, Message Interface Network150 then routes the message to the next destination (e.g., another service 350) on the

routing list. After passing through each of the intermediate destinations on the routing list, the message is then stored in queue 334 for recipient service 360. It should be noted that queues can also be associated with in-transit services 350. For simplicity, these queues are not shown in the illustrated embodiment of FIG. 3.

[0110] Application 372 in recipient service 360 can retrieve the message from queue 334 via message poll interface 328. In the poll mode, application 372 periodically issues a call to message poll interface 328 to ask for any waiting messages. If there are queued messages waiting for delivery to that service, then the one or more messages are returned in a reply. When making poll requests, a service can provide selectors on the messages to fetch. For example, a service can retrieve messages based upon the sender, message type, topic, etc.

[0111] In an alternative embodiment, message delivery is enabled through a push mode. In the push mode, the recipient would have its own server to which Message Interface Network150 can send messages. A service can specify a maximum number of tries, as well as the retry interval, for Message Interface Network150 to send the message to the service before aborting pushed delivery. A service to which Message Interface Network150 pushes messages can also optionally poll for messages.

[0112] In the push mode, a service can also specify a delivery window in which it will accept pushed messages. For example, a service might only accept messages between the hours of 1 AM and 2 AM.

[0113] As further illustrated in FIG. 3, a response message can be posted by recipient service 360 to message post interface 326. The message is then routed through one or more services 350 prior to being stored in message queue 332. The message can then be retrieved

by sending service 310 through message poll interface 322. As noted, the return path would not necessarily match the forward path.

[0114] In one embodiment, the sender of a message can specify a time by which the routing of a message must complete. If that expiration time passes before delivery of the message to its final destination, Message Interface Network150 will abort further routing of the message. In one embodiment, the final destination for a request message is the delivery of a response message back to the original request sender. In this embodiment, senders of response messages cannot specify an expiration since the response is considered part of the routing of the original request message. If a request message expires, Message Interface Network150 will return an error response back to the sender. If the sender does not specify an expiration time, then a default message expiration time (e.g., 48 hours) can be used. It should be noted that message expiration is not the same as document expiration. Document expiration is associated with a specific document and indicates how long the document's information is valid.

[0115] As part of the message delivery process, Message Interface Network150 logs all posted messages, including invalid messages. For each message, Message Interface Network150 logs relevant information to track the history of the message. Message Interface Network 150 also maintains a correlation between messages. That is, for request messages, Message Interface Network 150 associates the log of the response message(s) with the log of the request message.

[0116] In one embodiment, logged information can include the message header, the calculated route for the message (unless route calculation fails), the status of route validation, the size of the message at each stage of the route, and the routing history, GCENP004

including the status for each service along the message's route. The status values for each service depends on the role of the service.

[0117] Message Interface Network 150 correlates all messages for the same message transaction. That is, for request messages, Message Interface Network 150 associates the log of the response message(s) with the log of the request message. Similarly, if a message causes an error message, then Message Interface Network 150 associates the log of the generated error message with the log of the original message.

[0118] Fig. 4A is a diagrammatic representation of a communication exchange 400 between a plurality of services in accordance with one embodiment of the present invention. In general, any type of service may exchange messages with any other type of service through Message Interchange Network which operates to correlate such messages. As shown in Fig. 4A, a service 402 initially sends a request message 404 through Message Interchange Network 406 to Service 408. Service 408 sends a response message 410 back to Service 402 in reply to the request 404. Service 408 also sends a request message 411 to Service 412 in response to the original request message 404 from Service 402. Service 412 then sends a response message 414 back to Service 408 in reply to the request 411.

[0119] A service may represent any computer application or device accessible to other entities through an interface such as an application programming interface (API), user interface, or Internet web user interface, by any of a variety of protocols over a network within an entity or over the public Internet. A service may also be comprised of multiple methods or applications implemented on a single device or distributed across multiple devices.

[0120] In particular cases, a message may also be sent to a routing script, as described above. In a correlation implementation, a routing script can also be defined as a service. For example, an application type service may send a message to a routing script type service. Accordingly, routing programs can be session participants. In fact, one of the major uses of sessions may be that a session encompasses the set of calls into and from a routing script.

[0121] The messages in this communication exchange 400 are all related and are preferably correlated together. Response message 410 is in reply to request message 404, and hence, related to request message 404. Likewise, response message 414 is in reply to request 411, and hence, related to request message 411. These two pairs of messages may each be referred to as a call. As shown, request 404 and response 410 form a call, while request 411 and response 414 form a call.

[0122] In one embodiment, each message is given a Hop Identifier (ID) and each call is given a Call ID. As shown, request 404 and response 410 are given Call ID 416, and request 411 and response 414 are given Call ID 418. In this example, the messages belonging to Call ID 416 are also related to the messages belonging to Call ID 418 since the latter messages were initiated in response to the former messages. Hence, the messages associated with both Call ID 416 and 418 are given a Session ID 420. In sum, each message is given a Hop ID, each set of related messages exchanged between two services is given a Call ID, and each set of messages belonging to related calls are given a Session ID.

[0123] Fig. 4B illustrates a group of related messages within another example communication exchange 450 in accordance with one embodiment of the present invention. As shown, a first service 454 initiates a call between itself and service 452. In response to this initial call, service 452 then has a plurality of resulting calls between itself and a

plurality of services 456, 458, 460, and 462. As illustrated, a first invocation to a service may initiate the invocation of a second service. As described above, the invocations can be recursive and rather complex and are only limited to the computational and storage resources of the Message Interchange Network 406.

[0124] In this example, the messages of these related calls all have the same assigned unique Session ID. Similarly, each call has a unique Call ID, while each message within each call has a unique Hop ID. However, not all calls initiated by a particular service have to be part of the same session. For example, service 452 can choose that some of the resulting calls are not part of the same session. In one implementation, a call is a predefined set of messages (*e.g.*, as preconfigured in the Message Interchange Network), such as a notification message or a pair of request and response messages. A call may be predefined as any number and type of messages. A session in this implementation depends on how the particular services define sessions. That is, a service may be implemented to define its messages any suitable manner. For example, a service may send a first message with the intention of the message being part of a first session and send a second message with the intention of the message being part of a second session. In this example, the sending service may include a first Session ID in the first message and a second Session ID in the second message. Alternatively, the Message Interchange Network may predefine what constitutes a session. For instance, messages which are spawned from a service in response to a message received into such service may be defined as part of a same session.

[0125] In summary, a call may include any predefined number and type of messages. For instance, each call may include any number of request and response messages and/or any number of notification messages (*e.g.*, a response is not sent by the receiving service). Each

session may include any number of calls between any number of service pairs. While the definition of sessions is general, common uses of sessions include grouping calls comprising a conversation between services, comprising a transaction involving multiple services, or comprising a trace of routing a message through diverse services.

[0126] It should also be noted that a same message may be passed through any number of services or hops. That is, a particular message may result in multiple Hop ID's, one Hop ID for each hop of its route. A same message may also be modified by such services as it passes through such services. A message which enters a service, is modified by such service, and output from such service may be defined as two different messages: a first message which is received by such service and a second modified message which exits the service.

[0127] Any suitable technique may be used to correlate related messages. Fig. 4C is a flowchart illustrating a procedure 470 for correlating messages in accordance with one embodiment of the present invention. Initially, a message is received into Message Interchange Network 406 in operation 472. The message may be any form, such as a request, response, or notification message. Any of these three message types can also include error messages, alert messages, cancel messages, acknowledge messages, etc.

[0128] It is then determined whether this is the first message for the current session in operation 474. In other words, it is determined whether the received message is the start of a new session. This determination may be performed in any suitable manner. In one implementation, if the received message does not include a Session ID or a Call ID, then it is determined that it is the first message in the session (*e.g.*, that this is the start of a new session). The Message Interchange Network 406 then assigns a unique session ID for the

current session in operation 476. The Message Interchange Network 406 may also return the assigned Session ID to the sender of the message and embeds the Session ID in the received message before forwarding the message to the recipient service in operation 476. When the received message is not the first message in the current session, assignment of a new Session ID is skipped.

[0129] If the message is not the first message in the current session, a Session ID may still be returned back to the message sender in an alternative implementation, as well as storing information about the message. In either implementation, a new Session ID is not assigned, but rather the Session ID of the current session is used for correlating with the current message.

[0130] Messages which are sent by either the sender service or the receiver service after the first message will include the either the assigned Session ID or Call ID if they belong to the current session. A message which is in reply to a request message may simply include a Call ID which is the same as the Call ID of the corresponding request message. In this case, the response message does not include a Session ID and the Message Interchange Network 406 determines that the response message belongs to the session to which the Call ID is associated.

[0131] It may then be determined whether the received message is the first message in the current call in operation 478. In one example, if the message already includes a Call ID, then it is determined that the message is not the first message for such call. If the message does not include a Call ID, then it is determined that this is the first message for the current call and the Message Interchange Network 406 then assigns a Call ID in operation 480. The Call ID may also be returned to the message sender and embedded in the received message

in operation 480. If it is determined that this message is not the first message for the current call, then this operation 480 for assigning a new Call ID is skipped.

[0132] If the message is not the first message in the current call, a Call ID may still be returned back to the message sender in an alternative implementation, as well as storing information about the message. In either implementation, a new Call ID is not assigned, but rather the Call ID of the current call is used for correlating with the current message.

[0133] This Call ID assignment operation 480 is also performed after the Session ID assignment operation 476, where a determination is not made as to whether the message is the first message in the current call. When it is determined that the message is the first message for the current session, it is also assumed that it is the first message for the current call.

[0134] The Message Interchange Network 406 may then assign a Hop ID to the received message in operation 482. The Hop ID generally identifies the current message and its associated “hop” between its sender and receiver.

[0135] The Session ID, Call ID and Hop ID may be assigned to a message in any suitable order, besides the order illustrated in Fig. 4C. For particular protocols, the messages may not be able to embed a Session ID and/or a Call ID. For these protocols, the assigned Hop ID, Call ID, and Session ID and their corresponding message information may still be stored as described below.

[0136] Information regarding the received message, the current call, and the current session is then stored in operation 484. This information can be stored at the same time or individually for the received message, current call, and current session at various points in

the message handling process. No matter when the information is stored, it may be stored in any suitable one or more memory device(s).

[0137] In one implementation, the Hop ID, Call ID, and Session ID and their associated information regarding the message, call, and session are stored in a repository. Fig. 4D illustrates a Message Correlation Table 490 for correlating information regarding each message, a Call Correlation Table 492 for correlating information regarding each call, and a Session Correlation Table 494 for correlating information regarding each session in accordance with one embodiment of the present invention. The Tables 490, 492, and 494 may be stored in repository 240. The repository may be formed from any suitable number of databases stored on any suitable number of memories on any suitable number of computer systems.

[0138] As shown, the Message Correlation Table 490 includes one or more entries for each message. Each entry includes correlation information for each message. In this embodiment, each entry corresponds to a particular message and includes the particular message's Hop ID, message information regarding the particular message, the particular message's Call ID, and the particular message's Session ID. The message information may include any information relevant to the particular message, such as the message's sending service and receiving service, whether the message has completed its transmission (*e.g.*, to its intended destination service), and a reason or error log regarding why a message has failed to complete its hop, selected message content, the size of the message, the receiving and sending times for the message, the topic of the message, the status on processing steps taken on the message, the protocols used in receiving and sending the message, etc. The message's session ID may be assigned by the Message Interchange Network, embedded in

the message received into the Message Interchange Network, or inferred from the Call ID embedded in the message (e.g., the same session to which the Call ID is associated).

[0139] As shown, a message having Hop ID equal to “hop1” belongs to the same call as a message identified by Hop ID “hop2.” This same call is identified by Call ID “call1.” Additionally, another message having Hop ID “hop3” belongs to a call identified by Call ID “call2”, while another message having Hop ID “hop4” belongs to a call identified by Call ID “call3.” Also illustrated in the example of Fig. 4D, messages having Hop ID’s “hop1”, “hop2”, and “hop3” all belong to a same session identified by Session ID “session1.” In contrast, the message having Hop ID “hop4” belongs to a different session having a Session ID “session2.”

[0140] In an alternative embodiment, a service could define a message field that is to be used as a message identifier. For example, the profile information for a service could include an XPath expression designating a field in XML messages. In this embodiment, the Message Interchange Network would extract the value of the field from a received message and store the custom message identifier along with the assigned Call or Session ID. This alternative implementation then enables a service to query for correlation information using their own identifier, rather than necessarily the assigned Call or Session ID’s.

[0141] The Call Correlation Table 492 also includes one or more entries for each call. As shown, each entry includes a Call ID for a particular call and call information regarding the particular call. The call information for a particular call may include an indication as to whether the particular call is complete, a reason for the particular call not completing, the type of call, the receiving and sending times for the call, the sender and recipient services of the call, the status of policy evaluation for the call, the set of hops in the call, etc.

[0142] The Session Correlation Table 494 includes one or more entries for each session. Each entry for a particular session generally includes a Session ID for the particular session and session information regarding the particular session. The session information for a particular session may include any information relevant to a particular session such as, the calculated or executed route for messages sent within such particular session, the size of each message at each stage in the route, the identity and status of each service within the route paths of the particular session, whether a session has been completed, and reasons for a session failing to complete, the initiating time and completion time for the session, the set of calls in the session, etc.

[0143] The tables shown in figure 4D illustrate only one embodiment of storing correlation information. Any other suitable set of tables and relationships among tables may be used to store correlation information. For example, the message route of a session may not be stored directly in one table, but rather may be inferred from the set of calls and hops related to the session as stored in one or more other tables.

[0144] At any time during the correlation process, an entity may query the Message Interchange Network 406 for information which has been correlated together. For example, a service may request correlation information regarding a particular session. In this example, the Message Interchange Network 406 collects all entries within the Correlation Tables 490, 492, and 494 that relate to the requested session's Session ID. Likewise, an entity may request correlation information regarding a particular service, route, call, time period, message status, etc.

[0145] The Message Interchange Network 406 preferably determines whether a request service is authorized to receive a compilation of its requested correlation information. This

determination may be performed in any suitable manner. In one implementation, the requesting service is authorized when the requested information is in any way related to a service owned by the requesting service's organization. For example, a service may request correlation information regarding any messages sent by such requesting service, and any messages from calls or sessions belonging to the same call or session as the sent messages. In another implementation, the requesting service may request correlation information which is related to any one message which contains an identifier of the requesting service.

[0146] The correlation techniques of the present invention allow end points or services to correlate through an intermediary which handles ID assignments. Thus, the end points or services do not have to separately negotiate the use of an ID for their communication exchanges with every service. As a result, this approach greatly simplifies correlation of complex message interchanges over a computer network between diverse services and/or end point entities.

[0147] In summary, since an intermediary network (*e.g.*, message interchange network) handles identification of messages, calls, and sessions, as well as correlation of information related to same, users of such intermediary network do not have to implement any procedures for correlation. Since an intermediary is performing the correlation of messages sent between diverse and distributed services, any end point or service can also easily gain a clear picture as to what is happening with messages sent as a result of its interactions with such diverse services. Since the retained correlation information is performed by a single intermediary network, the correlation information is standardized and easily compiled into a clear summary of interactions results.

[0148] As thus described, Message Interface Network 150 enables flexible interaction between services. Several embodiments for provisioning or setting up services with Message Interface Network 150 and routing messages through Message Interface Network 150 between such provisioned services are further described in the above referenced Lev Brouk et al. U.S. patent application.

[0149] Generally, the correlation techniques of the present invention may be implemented on software and/or hardware. For example, it can be implemented in an operating system kernel, in a separate user process, in a library package bound into network applications, on a specially constructed machine, or on a network interface card. In a specific embodiment of this invention, the techniques of the present invention are implemented in software such as an operating system or in an application running on an operating system.

[0150] A software or software/hardware hybrid packet processing system of this invention is preferably implemented on a general-purpose programmable machine selectively activated or reconfigured by a computer program stored in memory. In one embodiment, portions of the correlation system (e.g., Message Interchange Network) may be implemented on a general-purpose network host machine such as a personal computer or workstation.

[0151] Referring now to Fig. 5 a computer system 500 suitable for implementing the present invention includes a master central processing unit (CPU) 502, one or more memory 504, input and output interfaces 506, and a bus 508 (e.g., a PCI bus). When acting under the control of appropriate software or firmware, the CPU 502 is responsible for implementing various portions of the techniques of the present invention. It preferably accomplishes all these functions under the control of software including an operating system and any appropriate applications software. CPU 502 may include one or more processors such as a

processor from the Intel family of microprocessors or the MIPS family of microprocessors.

In a specific embodiment, a memory 504 (such as non-volatile RAM and/or ROM) also forms part of CPU 502. However, there are many different ways in which memory could be coupled to the system. Memory block 504 may be used for a variety of purposes such as, for example, caching and/or storing data, programming instructions, etc.

[0152] The input and output interfaces 506 typically provide an interface to various I/O devices, such as mouse, keyboard, display, as well as providing a communication interface with other computer systems over a computer network. Among the communication interfaces that may be provided are Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like. In addition, various very high-speed interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM.

[0153] Although the system shown in Figure 5 is one specific computer system of the present invention, it is by no means the only system architecture on which the present invention can be implemented.

[0154] Regardless of system's configuration, it may employ one or more memories or memory modules (such as, for example, memory block 504) configured to store data, program instructions for the general-purpose network operations and/or the inventive techniques described herein. The program instructions may control the operation of an

operating system and/or one or more applications, for example. The memory or memories may also be configured to store information in repository 240, etc.

[0155] Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks and DVDs; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). The invention may also be embodied in a carrier wave travelling over an appropriate medium such as airwaves, optical lines, electric lines, etc. Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

[0156] Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. Therefore, the described embodiments should be taken as illustrative and not restrictive, and the invention should not be limited to the details given herein but should be defined by the following claims and their full scope of equivalents.